

 **CERT** *Ya!*

Microsoft

DP-300

Administering Relational Databases on Microsoft Azure

QUESTION & ANSWERS

QUESTION 1

Case Study	Number of Questions	Total Question
Case Study: 1	13	1 – 13
Case Study: 2	6	14 – 19
Case Study: 3	2	20 – 21
Case Study: 4	3	22 – 24
Case Study: 5	153	25 – 177
Total		177

Case Study: 1

Litware

Existing Environment

Network Environment

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

Database Environment

The sales department has the following database workload:

An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB

databases.

A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.

An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.

Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

Maintenance tasks must be automated.

The 30 new databases must scale automatically.

The use of an on-premises infrastructure must be minimized.

Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.

All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

Store encryption keys in Azure Key Vault.

Retain backups of the PII data for two months.

Encrypt the PII data at rest, in transit, and in use.

Use the principle of least privilege whenever possible. Authenticate database users by using Active Directory credentials.

Protect Azure SQL Database instances by using database-level firewall rules.

Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

Business Requirements

Litware identifies the following business requirements:

Meet an SLA of 99.99% availability for all Azure deployments.

Minimize downtime during the migration of the SERVER1 databases.

Use the Azure Hybrid Use Benefits when migrating workloads to Azure.

Once all requirements are met, minimize costs whenever possible.

Question No. 1

You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

What should you include in the migration plan? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Azure Database Migration Service pricing tier:

	▼
Standard 2-vCore	
Standard 4-vCore	
Premium 4-vCore	

Required Azure resource:

	▼
A virtual network that has service endpoints	
A VPN gateway	
An Azure Logic app	

Correct Answer:

Azure Database Migration Service pricing tier:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">Standard 2-vCore</td></tr><tr><td colspan="2">Standard 4-vCore</td></tr><tr><td colspan="2">Premium 4-vCore</td></tr></table>		▼	Standard 2-vCore		Standard 4-vCore		Premium 4-vCore	
	▼								
Standard 2-vCore									
Standard 4-vCore									
Premium 4-vCore									
Required Azure resource:	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">A virtual network that has service endpoints</td></tr><tr><td colspan="2">A VPN gateway</td></tr><tr><td colspan="2">An Azure Logic app</td></tr></table>		▼	A virtual network that has service endpoints		A VPN gateway		An Azure Logic app	
	▼								
A virtual network that has service endpoints									
A VPN gateway									
An Azure Logic app									

Explanation/Reference:

Explanation:

<https://azure.microsoft.com/pricing/details/database-migration/>

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online>

QUESTION 2

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an

Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You run the Remove-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. You run the Restore-AzSqlDatabase PowerShell cmdlet for Database1 on Server2. Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Explanation/Reference:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

<https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

QUESTION 3

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state. What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

Correct Answer: B

Explanation/Reference:

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the maximum size limit between storage size checks, it transitions to read-only mode.

Incorrect Answers:

C: Statistics Collection Interval: Defines the level of granularity for the collected runtime statistic,

expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues. Keep in mind that the value directly affects the size of Query Store data. <https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

QUESTION 4

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
(
    EmployeeID          INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName           VARCHAR(100) NOT NULL,
    LastName            VARCHAR(100) NOT NULL,
    Title               VARCHAR(100) NULL,
    LastHireDate        DATETIME NULL,
    StreetAddress1      VARCHAR(500) NOT NULL,
    StreetAddress2      VARCHAR(500) NOT NULL,
    StreetAddress3      VARCHAR(500) NOT NULL,
    City                VARCHAR(200) NOT NULL,
    StateName           VARCHAR(20) NOT NULL,
    Salary              VARCHAR(20) NULL,
    PhoneNumber         VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Change Salary to the money data type.
- B. Change PhoneNumber to the float data type.
- C. Change LastHireDate to the datetime2(7) data type.
- D. Change PhoneNumber to the bigint data type.
- E. Change LastHireDate to the date data type.

Correct Answer: A,E

QUESTION 5

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You schedule an Azure Databricks job that executes an R notebook, and then inserts the data into the data warehouse.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Explanation/Reference:

Must use an Azure Data Factory, not an Azure Databricks job.

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

QUESTION 6

You are designing an enterprise data warehouse in Azure Synapse Analytics that will store website traffic analytics in a star schema.

You plan to have a fact table for website visits. The table will be approximately 5 GB.

You need to recommend which distribution type and index type to use for the table. The solution must provide the fastest query performance.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Distribution:

	▼
Hash	
Round robin	
Replicated	

Index:

	▼
Clustered columnstore	
Clustered	
Nonclustered	

Distribution: ▼
Hash
Round robin
Replicated

Index: ▼
Clustered columnstore
Clustered
Nonclustered

Correct Answer:

Explanation/Reference:

Explanation:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-index>

QUESTION 7

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1.

DF1 contains a linked service.

You have an Azure Key vault named vault1 that contains an encryption key named key1.

You need to encrypt df1 by using key1.

What should you do first?

- A. Disable purge protection on vault1.
- B. Remove the linked service from df1.
- C. Create a self-hosted integration runtime.
- D. Disable soft delete on vault1.

Correct Answer: B

Explanation/Reference:

A customer-managed key can only be configured on an empty data Factory. The data factory can't contain any resources such as linked services, pipelines and data flows. It is recommended to enable customer-managed key right after factory creation.

Note: Azure Data Factory encrypts data at rest, including entity definitions and any data cached while runs are in progress. By default, data is encrypted with a randomly generated Microsoft-managed key that is uniquely assigned to your data factory.

<https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key>

QUESTION 8

You need to implement the surrogate key for the retail store table. The solution must meet the sales transaction dataset requirements.

What should you create?

- A. a table that has a FOREIGN KEY constraint
- B. a table that has an IDENTITY property
- C. a user-defined SEQUENCE object
- D. a system-versioned temporal table

Correct Answer: B

Explanation/Reference:

Scenario: Contoso requirements for the sales transaction dataset include:

Implement a surrogate key to account for changes to the retail store addresses.

A surrogate key on a table is a column with a unique identifier for each row. The key is not generated from the

table data. Data modelers like to create surrogate keys on their tables when they design data warehouse models. You can use the IDENTITY property to achieve this goal simply and effectively without affecting load performance.

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesidentity>

QUESTION 9

You have an Azure SQL managed instance.

You need to gather the last execution of a query plan and its runtime statistics. The solution must minimize the impact on currently running queries. What should you do?

- A. Generate an estimated execution plan.
- B. Generate an actual execution plan.
- C. Run sys.dm_exec_query_plan_scans.
- D. Generate Live Query Statistics.

Correct Answer: C

Explanation/Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15>

QUESTION 10

You need to implement a solution to notify the administrators. The solution must meet the monitoring requirements.

What should you do?

- A. Create an Azure Monitor alert rule that has a static threshold and assign the alert rule to an action group.
- B. Add a diagnostic setting that logs QueryStoreRuntimeStatistics and streams to an Azure event hub.
- C. Add a diagnostic setting that logs Timeouts and streams to an Azure event hub.
- D. <https://www.mytestengine.com/home/index?code=DP-300>

Correct Answer: D

Explanation/Reference:

<https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/>

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.

QUESTION 11

You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database.

Which database option should set?

- A. PARAMETERIZATION to FORCED

- B. PARAMETERIZATION to SIMPLE
- C. Delayed Durability to Forced
- D. READ_COMMITTED_SNAPSHOT to ON

Correct Answer: D

Explanation/Reference:

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON.

The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Incorrect Answers:

A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may choose to parameterize the queries. This means that any literal values that are contained in a query are substituted with parameters. This process is referred to as simple parameterization. When SIMPLE parameterization is in effect, you cannot control which queries are parameterized and which queries are not.

B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION database option to FORCED. This process is referred to as forced parameterization.

C: Delayed transaction durability is accomplished using asynchronous log writes to disk. Transaction log records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place. Delayed transaction durability reduces both latency and contention within the system. Some of the cases in which you could benefit from using delayed transaction durability are:

You can tolerate some data loss.


You are experiencing a bottleneck on transaction log writes.

Your workloads have a high contention rate.


<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql>

QUESTION 12

You have an Azure data factory that has two pipelines named PipelineA and PipelineB. PipelineA has four activities as shown in the following exhibit.



PipelineB has two activities as shown in the following exhibit.



You create an alert for the data factory that uses Failed pipeline runs metrics for both pipelines and all failure types. The metric has the following settings:

Operator: Greater than

Aggregation type: Total

Threshold value: 2

Aggregation granularity (Period): 5 minutes

Frequency of evaluation: Every 5 minutes

Data Factory monitoring records the failures shown in the following table.

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification was sent after the failure of Activity3 in PipelineA.	<input type="radio"/>	<input checked="" type="radio"/>
An alert notification was sent after the failure of Activity1 in PipelineB.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation/Reference:

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview>

QUESTION 13

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.
What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the dbcc dbreindex command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Correct Answer: C

Explanation/Reference:

Explanation:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview>

QUESTION 14

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1. The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.


From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.



Correct Answer:

Set up the Active Directory Admin for AzSQL1.
Connect to DB1 by using the server administrator.
Create a user by using the FROM EXTERNAL PROVIDER clause.



Set up the Active Directory Admin for AzSQL1.
Connect to DB1 by using the server administrator.
Create a user by using the FROM EXTERNAL PROVIDER clause.

Explanation/Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql>

QUESTION 15

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:

The first six numerals of each customer's phone number must be masked.

The last four digits of each customer's phone number must be visible.

Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Exposed Prefix:

	▼
0	
1	
3	
5	

Padding String:

	▼
x	
xxxxxx	
xxx-xxx	
xxx-xxx-	
x[3]-x[3]	

Exposed Suffix:

	▼
0	
1	
3	
5	

Exposed Prefix: ▼

0
1
3
5

Padding String: ▼

x
xxxxxx
xxx-xxx
xxx-xxx-
x[3]-x[3]

Exposed Suffix: ▼

0
1
3
5

Correct Answer:

Explanation/Reference:

Explanation:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>